arXiv:1412.7334v1 [stat.AP] 23 Dec 2014

A hidden Markov approach to disability insurance

Boualem Djehiche* Björn Löfdahl †

December 24, 2014

Abstract

Point and interval estimation of future disability inception and recovery rates are predominantly carried out by combining generalized linear models (GLM) with time series forecasting techniques into a two-step method involving parameter estimation from historical data and subsequent calibration of a time series model. This approach may in fact lead to both conceptual and numerical problems since any time trend components of the model are incoherently treated as both model parameters and realizations of a stochastic process. We suggest that this general two-step approach can be improved in the following way: First, we assume a stochastic process form for the time trend component. The corresponding transition densities are then incorporated into the likelihood, and the model parameters are estimated using the Expectation-Maximization algorithm. We illustrate the modelling procedure by fitting the model to Swedish disability claims data.

Keywords: Disability insurance, Hidden Markov model, Maximum Likelihood, Expectation-Maximization.

1 Introduction

To determine premiums and reserves associated with health and disability insurance policies, the insurer needs predictions of the future rates of disability inception and recovery. While earlier research provides a solid base for disability modelling, new studies are required as the field is in constant change due to policy reforms and amendments to the existing regulations. For instance, the Swedish government launched major reforms of the national sickness insurance system in 2008, changing the rules for obtaining benefits from the Social Insurance Agency. This reform has been of major importance to the reduction in sickness absence. As noted by Aro *et al.* [1], research of sickness and disability on Swedish data undertaken before 2008 may no longer provide an accurate description of the disability dynamics after the reform. As of October 2014, the Swedish government has suggested that the reforms of 2008 should essentially be reversed, a change that would require a drastic increase in premiums and reserves. This proposal highlights the need to study the calendar time dynamics of disability.

*Department of Mathematics, KTH Royal Institute of Technology, Sweden, boualem@kth.se.
†Department of Mathematics, KTH Royal Institute of Technology, Sweden, bjornlg@kth.se.

A popular approach to estimating disability inception and recovery rates is based on the generalized linear models framework. Renshaw and Haberman [12] model recovery, mortality and inception time trends in permanent health insurance using Poisson regression. Christiansen *et al.* [4] model recovery, mortality and inception using the functional data approach of Hyndman and Ullah [10]. Aro *et al.* [1] propose logistic regression models for disability inception and termination. Similar methods are used for modelling population mortality, where perhaps the most well known is the Lee-Carter model [11] and its extensions, including the Poisson log-bilinear model of Brouhns *et al.* [3].

For the purpose of obtaining point and interval estimates of future disability or mortality rates, it is customary to combine GLMs with time series forecasting techniques into the following two-step method: In the first step, the parameters of the GLM are estimated from historical data. In the second step, a time trend component ν_t is assumed to follow a time series model, where a popular choice is the random walk with drift, and the parameters of the model are fitted to the estimated values of $\{\nu_t\}$. Prediction or simulation of future transition rates are obtained by prediction or simulation from the time series model for ν_t. This two-step approach provides an easy way of fitting the model to data and simulating future outcomes, and it has been employed by e.g. Brouhns *et al.* [3], Christiansen *et al.* [4], Djehiche and Löfdahl [7] and others.

An issue with the two-step approach is that at first, $\{\nu_t\}$ are considered parameters to be estimated. After estimating them, the assumption is altered so that ν is treated as a stochastic process. This may lead to both conceptual and numerical problems. In particular, the volatility of ν tends to be overestimated, which may have significant impact on pricing and risk management of insurance products. Qualitatively, this stems from the fact that yearly variations in the parameter values are caused by variations in the underlying process (systematic variation) as well as variations in the underlying population (idiosyncratic variation). The two-step approach makes no distinction between idiosyncratic and systematic variations.

The incoherence of the two-step approach has previously been pointed out by Czado *et al.* [5]. They propose to avoid this deficiency by integrating both steps into a Bayesian model, where the yearly values of the process corresponding to ν are all treated as random variables with given prior densities. The model parameters are then estimated using the Gibbs sampler and Metropolis-Hastings algorithm.

We suggest an alternative solution to this problem that does not rely on Bayesian statistics: First, we assume a stochastic process form for the time component ν. Then, we fit the model over all time periods simultaneously, incorporating the transition densities of ν into the likelihood. Maximization of the likelihood cannot, however, be carried out directly, since now we no longer consider $\{\nu_t\}$ to be parameters. Instead, ν is treated as an unobservable stochastic process, so that the model is formulated as a Hidden Markov model (HMM). We then proceed using the Expectation-Maximization (EM) algorithm of Dempster *et al.* [6]. To the best of our knowledge, this type of EM-algorithm has not been used before for the purpose of estimating disability rates. It is, however, well known in other fields, such as finance [8].

As a starting point for the improved model, we consider the disability model from Aro *et al.* [1]: Let $E_{x,t}$ be the number of healthy individuals aged x at the beginning of time period t in a given disability insurance scheme. We denote by $D_{x,t}$ the number of individuals falling ill amongst the $E_{x,t}$ insured healthy individuals during time interval $[t, t+1)$. The authors assume that the conditional distribution of $D_{x,t}$ given $E_{x,t}$ is binomial:

$$D_{x,t} \sim \mathrm{Bin}(E_{x,t}, p_{x,t}), \tag{1}$$

where $p_{x,t}$ is the probability that an x-year-old individual randomly selected at t falls ill during $[t, t+1)$. Further, they suggest to model the logistic disability inception probabilities by

$$\operatorname{logit} p_{x,t} = \sum_{i=1}^{m} \nu_t^i \phi^i(x), \tag{2}$$

where ϕ^i are user-defined basis functions, and ν_t^i are risk factors to be estimated from data. The authors also propose a straightforward extension of this model to disability termination modelling.

The historical values of the risk factors $\nu_t = (\nu_t^1, \ldots, \nu_t^m)$ can be easily obtained by maximum likelihood estimation as follows. Given the historical values of $D_{x,t}$ and $E_{x,t}$, the log-likelihood function for yearly values of ν_t can be written using (1) and (2) as

$$l_t(\nu_t; D_{\cdot,t}) = \sum_{x \in X} \Big[D_{x,t} \sum_{i=1}^{m} \nu_t^i \phi^i(x) - E_{x,t} \log \big(1 + \exp\big\{ \sum_{i=1}^{m} \nu_t^i \phi^i(x) \big\}\big) \Big]. \tag{3}$$

Aro and Pennanen [2] show that if the basis functions are linearly independent, the yearly log-likelihood $l_t(\cdot; D_{\cdot,t})$ is strictly concave. Hence, maximizing $l_t(\cdot; D_{\cdot,t})$ over $\nu_t \in \mathbb{R}^m$ using numerical methods gives a unique estimate of the vector ν_t for each t.

In this paper, we propose instead to treat ν as a hidden Markov process with transition densities parameterized by θ, say, and use the Expectation-Maximization algorithm as follows: Given a parameter estimate θ^k, integrate the complete data log-likelihood $l(\theta; D_{\cdot,1:n}, \nu_{1:n})$ with respect to the distribution of $\nu_{1:n} := (\nu_1, \ldots, \nu_n)$ conditional on the observations $D_{\cdot,1:n} := (D_{\cdot,1}, \ldots, D_{\cdot,n})$, e.g. let

$$Q(\theta|\theta^k) = E^{\theta^k}[l(\theta; D_{\cdot,1:n}, \nu_{1:n})|D_{\cdot,1:n}]. \tag{4}$$

Next, we maximize Q w.r.t. θ to obtain

$$\theta^{k+1} = \operatorname*{argmax}_{\theta} Q(\theta|\theta^k). \tag{5}$$

Iterating over expectation and maximization steps, the output of the EM-algorithm is a sequence $\{\theta^k\}$ of parameter estimates. Under technical conditions that are usually hard to verify, the sequence $\{\theta^k\}$ will eventually converge to a stationary point θ^* with $L(\theta^*) = L^*$ being the corresponding stationary point of the log-likelihood function. If the likelihood function is also unimodal, $\{\theta^k\}$ will converge to $\theta^* = \operatorname{argmax}_\theta L(\theta)$. See Wu [14] for details.

The outline of this paper is as follows. In Section 2, we propose a model for disability inception rates and show how the model parameters can be estimated using the EM algorithm. Section 3 illustrates the modelling procedure by fitting the model to disability claims data from the Swedish insurance company Folksam. In Section 4, we propose a version of the model for the estimation of disability termination rates. In Section 5, we fit the termination model to disability claims data from Folksam.

2 Disability inception model

Let $E_{x,t}$ be the number of healthy individuals aged x at the beginning of time period t in a given disability insurance scheme. We denote by $D_{x,t}$ the number of individuals falling ill amongst the $E_{x,t}$ insured healthy individuals during time interval $[t, t+$

1). Further, let ν be an m-dimensional Brownian motion starting at ν_0 with drift μ and Cholesky matrix A. This choice of ν corresponds directly to the frequently used ARIMA(0,1,0) random walk. We assume that the conditional distribution of $D_{x,t}$ given $E_{x,t}$ and ν_t is binomial:

$$D_{x,t} \sim \text{Bin}(E_{x,t}, p_{x,t}), \tag{6}$$

where $p_{x,t}$ given by

$$p_{x,t} := \frac{1}{1 + e^{-g(x,\nu_t)}}, \tag{7}$$

is the probability that an individual randomly selected at t falls ill during $[t, t + 1)$. Here, the selection of $g : \mathbb{R}^+ \times \mathbb{R}^m \mapsto \mathbb{R}$ is a model choice. We adopt the basis function approach from [1], and choose a g of the form

$$g(x, \nu_t) = \sum_{i=1}^{m} \nu_t^i \phi^i(x), \tag{8}$$

where $\phi^i,\ i = 1, \ldots, m$, are user-defined basis function.

Now, assume that we observe $D_{x,1:n} := (D_{x,1}, \ldots, D_{x,n})$ and $E_{x,1:n} := (E_{x,1}, \ldots, E_{x,n})$, for x from a given set X of ages. Let $\theta = (\mu, A, \nu_0)$. Then, the complete data log-likelihood is given by

$$l(\theta; D_{\cdot,1:n}, \nu_{1:n}) = \sum_{t=1}^{n} \Big[l_t(\nu_t; D_{\cdot,t}) + \log f_{\nu_t|\nu_{t-1}}(\theta) + c_t \Big], \tag{9}$$

where f is the density of ν_t given ν_{t-1}, and c_t is a constant. From the Brownian motion assumption, we have

$$\begin{aligned} \log f_{\nu_t|\nu_{t-1}}(\theta) = &- (\nu_t - \nu_{t-1} - \mu)^T (AA^T)^{-1} (\nu_t - \nu_{t-1} - \mu) \\ &- \frac{1}{2} \log(\det(AA^T)). \end{aligned} \tag{10}$$

This is a direct extension of the model from [1], in that instead of fitting each time period separately, we consider all time periods simultaneously by summation of the log-likelihood over all time periods. In addition, we include a term corresponding to the density of ν_t given ν_{t-1}.

The following proposition is useful for obtaining point estimates of and confidence intervals for ν_t.

Proposition 1 *The filter density functions $\phi_{\nu_t}(\cdot) = f_{\nu_t|D_{\cdot,1:t}}(\cdot)$ are log-concave on $\mathbb{R}^m$.*

Proof. Using Bayes' theorem, the law of total probability and the Markov property of ν, we can write $f_{\nu_t|D_{\cdot,1:t}}(x_t)$ as

$$f_{\nu_t|D_{\cdot,1:t}}(x_t) \propto p_{D_{\cdot,t}|\nu_t}(x_t) I(x_t), \tag{11}$$

where

$$I(x_t) = \int f_{\nu_t|\nu_{t-1}}(x_t - x_{t-1}) \prod_{k=1}^{t-1} p_{D_{\cdot,k}|\nu_k}(x_k) f_{\nu_k|\nu_{k-1}}(x_k - x_{k-1}) dx_{1:t-1}. \tag{12}$$

Since $\log p_{D_{\cdot,t}|\nu_t}(x_t) = l_t(x_t; D_t)$ is concave, it remains to show that $I(x_t)$ is log-concave. It is well known that the densities $f_{\nu_k|\nu_{k-1}}(x_k - x_{k-1})$ are log-concave on $\mathbb{R}^m \times \mathbb{R}^m$. Hence, the integrand $h(x_t, x_{1:t-1})$ defined by

$$h(x_t, x_{1:t-1}) = f_{\nu_t|\nu_{t-1}}(x_t - x_{t-1}) \prod_{k=1}^{t-1} p_{D_{\cdot,k}|\nu_k}(x_k) f_{\nu_k|\nu_{k-1}}(x_k - x_{k-1}) \qquad (13)$$

is log-concave on $\mathbb{R}^m \times \mathbb{R}^{m(t-1)}$. From [9, Corollary 2], log-concavity of h directly implies log-concavity of $I(x_t)$ on $\mathbb{R}^m$. □

It follows from Proposition 1 that the filter distributions are unimodal. This is a type of identification attribute of the model: Estimating the historical values of ν using the filter densities ϕ_{ν_t} admits identification of $\nu_t,\ t = 1, \ldots, n$, by their respective modes. Confidence intervals for historical values of ν_t can be obtained directly as quantiles of ϕ_{ν_t}.

Unfortunately, the filter distributions cannot be calculated directly. However, it is relatively easy to sample from them using particle filter methods, given that we have estimates of θ. By sampling from the filter distributions we can then obtain an updated estimate of θ using the Expectation-Maximization algorithm, with ν treated as a hidden Markov process. The choice of θ^0 is important since we have not been able to show that the log-likelihood function is unimodal. We suggest choosing θ^0 by fitting the model from Aro *et al.* for $t = 1, \ldots, n$ and estimating θ^0 from the time series of estimated values of ν_t. This procedure should yield a good start guess for the parameters ν_0 and μ, while the start guess for the standard deviations given by $\sqrt{\mathrm{diag}(AA^T)}$ should be overestimated.

Integrating the log-likelihood and discarding all terms that do not depend on θ, we obtain

$$Q(\theta|\theta^k) = \sum_{t=1}^{n} \Big[-\frac{1}{2} E^{\theta^k}[(\nu_t - \nu_{t-1} - \mu)^T (AA^T)^{-1} (\nu_t - \nu_{t-1} - \mu) | D_{\cdot,1:n}] - \frac{1}{2} \log(\det(AA^T)) \Big]. \qquad (14)$$

In order to maximize Q w.r.t. θ, we need to evaluate the conditional expectations appearing in (14). This is not a trivial problem, since it is required to determine the density of $\nu_t - \nu_{t-1}$ conditional on $D_{\cdot,1:n}$ using Bayes' theorem. However, there exist numerical techniques that allow for efficient evaluation of the expectations, including particle filter methods. Further, we need to write Q on a form that allows for easy maximization. We start with the latter task.

2.1 Maximization

Simple but tedious linear algebra yields the following expression for Q:

$$Q(\theta|\theta^k) = -\frac{n}{2} \log(\det(AA^T)) - \frac{1}{2} tr((AA^T)^{-1} C^T), \qquad (15)$$

where

$$\begin{aligned} C_{ij} =& S_{ij} - \mu_i S_j - \mu_j S_i + n\mu_i\mu_j + E_{ij} - \nu_0^i E_j - \nu_0^j E_i + \nu_0^i\nu_0^j \\ & - \mu_i(E_j - \nu_0^j) - \mu_j(E_i - \nu_0^i), \end{aligned} \tag{16}$$

$$S_{ij} = \sum_{t=2}^{n} E^{\theta^k}[(\nu_t^i - \nu_{t-1}^i)(\nu_t^j - \nu_{t-1}^j)|D_{\cdot,1:n}], \tag{17}$$

$$S_i = \sum_{t=2}^{n} E^{\theta^k}[(\nu_t^i - \nu_{t-1}^i)|D_{\cdot,1:n}], \tag{18}$$

$$E_{ij} = E^{\theta^k}[\nu_1^i\nu_1^j|D_{\cdot,1:n}], \tag{19}$$

$$E_i = E^{\theta^k}[\nu_1^i|D_{\cdot,1:n}]. \tag{20}$$

Given $S_{ij},\ S_i,\ E_{ij}$ and E_i for $i, j = 1, \ldots, m$, it is a simple matter to maximize (15) in the following way: First, taking derivatives, the optimal μ and ν_0 are given by

$$\mu_i = \frac{1}{n-1} S_i, \tag{21}$$

$$\nu_0^i = E_i - \frac{1}{n-1} S_i. \tag{22}$$

Substituting μ and ν_0 back into C_{ij} yields

$$C_{ij} = S_{ij} + E_{ij} - \frac{1}{n-1} S_i S_j - E_i E_j. \tag{23}$$

Now, since C is no longer a function of θ, it suffices to consider the mapping $A \mapsto \tilde{Q}(A)$ defined by

$$\begin{aligned} \tilde{Q}(A) =& -\frac{n}{2}\log(\det(AA^T)) - \frac{1}{2} tr((AA^T)^{-1}C^T) \\ =& -\frac{n}{2}\log(\det(AA^T)) - \frac{n}{2} tr((AA^T)^{-1}\bar{C}^T), \end{aligned} \tag{24}$$

where $\bar{C}_{ij} = \frac{1}{n} C_{ij}$. It is well known that (24) obtains its maximum value at

$$AA^T = \bar{C}, \tag{25}$$

provided that $\bar{C}$ is positive definite. Occasionally, due to Monte Carlo error, it may happen that $\bar{C}$ is not positive definite. This can be remedied in several ways, we may for example attempt to maximize (24) numerically. Another option is to resample and perform the E-step anew, and attempt the M-step once more using the updated expectations.

2.2 Expectation

We now turn towards the task of computing the conditional expectations. The conditional expectations $S_{ij},\ S_i,\ E_{ij}$ and E_i are of a form suitable to the particle-based rapid incremental smoother, or PaRIS, algorithm due to Westerborn and Olsson [13].

A particle filter is a necessary requirement for implementing the PaRIS algorithm, and for this purpose we choose to implement a simple bootstrap particle filter. The filter distributions ϕ_{ν_t}, that is, for each t, the distribution of ν_t conditional on $D_{\cdot,1:t}$,

are estimated in the following way: Given a sample of ν_{t-1}, we first sample N particles of ν_t from $f_{\nu_t|\nu_{t-1}}(\theta^k)$ to obtain $z_k = (z_k^i)_{i=1,\dots,m},\ k = 1,\dots,N$. Each particle z_k is then given the weight $w_k \propto \exp\{l_t(z_k; D_{\cdot,1:t})\}$, and the filter probability mass function is estimated by $\hat{\phi}_{\nu_t}(z_k) = w_k$. Finally, we bootstrap from $z_k,\ k = 1,\dots,N$ with probabilities $w_k,\ k = 1,\dots,N$, to obtain a sample of ν_t, and repeat the procedure until $t = n$. We estimate the yearly values of ν_t for $t = 1,\dots,n$ by

$$\widehat{\nu}_t^i = \sum_{k=1}^{N} w_k^i z_k^i,\ i = 1,\dots,m. \tag{26}$$

Confidence intervals for $\nu_t^i,\ i = 1,\dots,m$, are obtained by calculating the empirical quantiles based on $\hat{\phi}_{\nu_t}(\cdot)$. Finally, the expectations $S_{ij},\ S_i,\ E_{ij}$ and E_i are estimated using the PaRIS algorithm as outlined in [13].

Note that this separation of expectation and maximization is possible due to the model specification: Since the basis functions are chosen *a priori* and are not themselves to be estimated, all terms of Q that depend on both θ and ν are product terms. No other non-multiplicative dependencies are present. This allows us to write Q of the form (15), which allows for separating the expectation and maximization steps as required.

Consider the case where the model specification was written so that the basis functions were also to be estimated from data. For example, we may consider the Lee-Carter type model from [3] for the force of mortality $q_{x,t}$:

$$q_{x,t} := e^{\alpha_x + \beta_x \kappa_t}, \tag{27}$$

where α_x and β_x are to be estimated along with κ_t. Then, Q will contain terms $T_{x,t}$ of the form

$$T_{x,t} = E_{x,t} e^{\alpha_x} E^{\theta^k}[e^{\beta_x \kappa_t} | D_{\cdot,1:t}]. \tag{28}$$

In our approach, the conditional expectation $E^{\theta^k}[e^{\beta_x \kappa_t} | D_{\cdot,1:t}]$ can only be estimated for a fixed β_x. Hence, we cannot feasibly implement this version of the EM-algorithm for the model specified by (27), except by estimating this quantity over a range of values for β_x and using interpolation and extrapolation over this range in the M-step. It is, however, possible to fit a model of the type (8) to mortality data using the techniques of this Section. For a discussion on how to choose the basis functions $\{\phi^i\}$ for mortality modelling we refer to Aro and Pennanen [2].

3 Fitting Swedish disability inception rates

In this section, we implement the EM-algorithm from Sections 2.1-2.2 for the disability inception model from Section 2, and fit it to population data from Folksam.

3.1 Two-factor model

We implement the model from Section 2 with basis functions given by

$$\phi^1(x) = \frac{64 - x}{39} \quad \text{and} \quad \phi^2(x) = \frac{x - 25}{39},$$

for $x \in [25, 64]$. The linear combination $\sum_{i=1}^{2} \nu_t^i \phi^i(x)$ is also linear. Note that the same linear form for the curve of logit $p_{\cdot,t}$ could have been obtained using any two

linearly independent linear basis functions. However, this particular choice ensures a certain natural interpretation of the stochastic process ν. Namely, for every t,

$$\operatorname{logit} p_{25,t} = \nu_t^1 \phi^1(25) + \nu_t^2 \phi^2(25) = \nu_t^1,$$

and, similarly, $\operatorname{logit} p_{64,t} = v_t^2$. Hence, two components of ν represent the logit disability inception probabilities of ages 25 and 64, respectively.

The EM-algorithm stabilizes to within Monte Carlo error after about 180 iterations. We run it for 20 more iterations and estimate θ as the average over the 20 last iterations. The value of $Q(\theta^k|\theta^{k-1})$ for $k = 1, \ldots, 200$ is presented in Figure 1. The estimated inception probabilities from the Hidden Markov model (HMM) for the years 2000-2011 are displayed in Figures 2-3. For reference, they are compared to the estimations from [1], (hereafter referred to as the multi-period model). Note that, due to confidentiality, the actual values of the estimates are not reported. Figures 4-5 display the estimated filter densities for ν^1 and ν^2, respectively. Indeed, as inferred from Proposition 1, the estimated filter distributions are for the most part unimodal. They also seem to be symmetric, which makes estimation of the yearly values of ν^1 and ν^2 from their corresponding mean values or modes equivalent. Table 1 displays the estimated drift and volatility parameters from the HMM as a fraction of the corresponding estimates from the multi-period model.

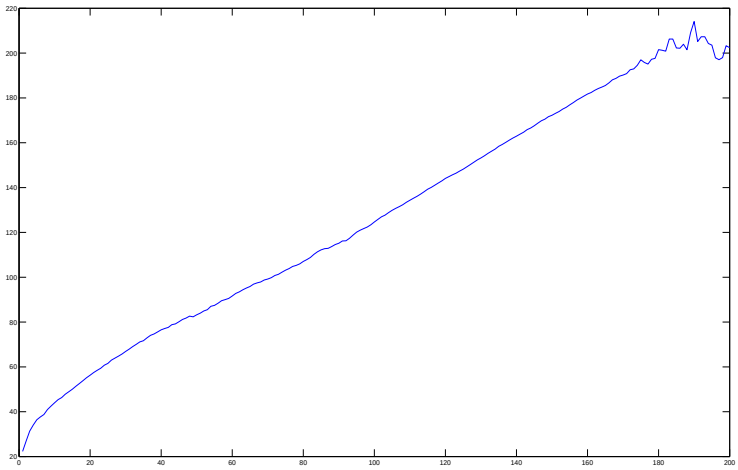

Figure 1: The value of $Q(\theta^k|\theta^{k-1})$ for $k = 1, \ldots, 200$.

Table 1: Relative difference of the estimated drift and volatility parameters between the two models.

	μ	σ
ν^1	0.92	0.48
ν^2	0.93	0.23

The HMM seems to provide estimates of $\nu_{1:n}$ and μ that are quite close to the estimates from the multi-period model, where the model is fitted to data for each time

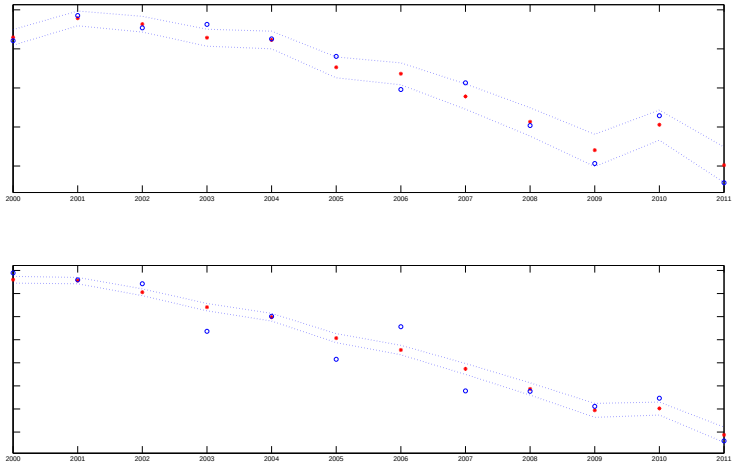

Figure 2: Estimates of $\nu_{1:n}$ (stars) with confidence bands (dashed). Estimates from [1] (circles) for comparison.

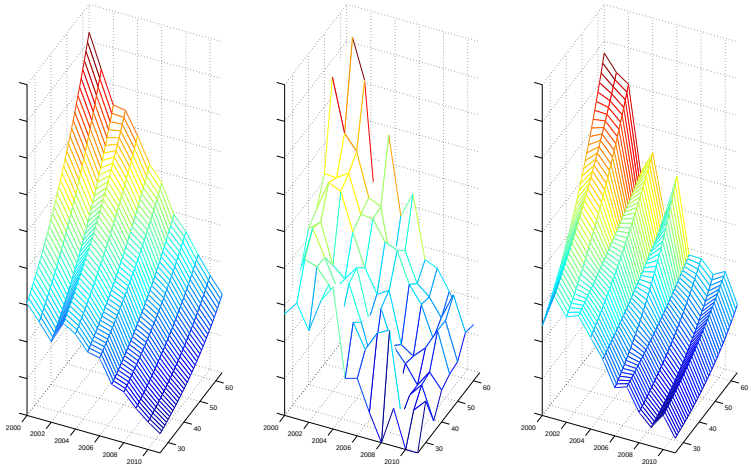

Figure 3: Left: Estimates of $p_{25:64,1:n}$. Middle: Raw data $D_{25:64,1:n}/E_{25:64,1:n}$. Right: Estimates from [1].

period separately. The goodness of fit as measured by the yearly log-likelihood values $l_t(\nu_t; D_{\cdot,t})$ is worse for the HMM, but this is to be expected: Calibration on only one time period will yield a better fit to the data corresponding to that particular time period compared to simultaneous calibration over many time periods, unless the estimates of ν_t turn out identical. The HMM provides smoothing across time due to the fact that the yearly estimates are essentially parameterized by a stochastic process.

The purpose of calibrating a model is usually not to obtain the best possible fit

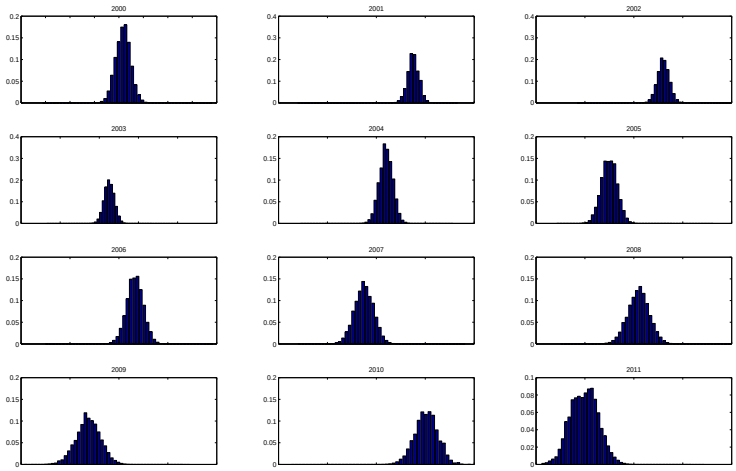

Figure 4: Estimated filter densities for ν^1.

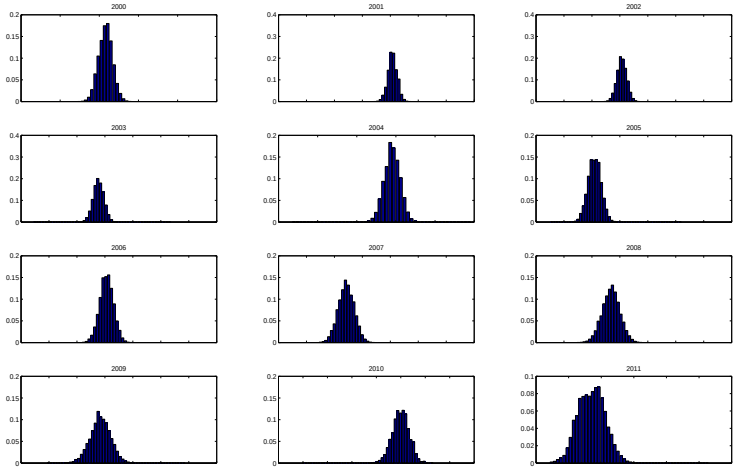

Figure 5: Estimated filter densities for ν^2.

to historical data. Rather, the objective should be to obtain the best fit of the law of future disability rates. The estimated diffusion coefficients from the HMM are considerably reduced compared to the multi-period model, which has a significant impact on the generation of future disability rates, prices, risks and capital charges. This is both reasonable expected. Qualitatively, this stems from the fact that yearly variations in parameter values are caused by variations in the underlying process (systematic variation) as well as variations in the underlying population (idiosyncratic variation). The multi-period model makes no distinction between idiosyncratic and systematic variations. The HMM, on the other hand, enables us to distinguish between the two, and in

effect the yearly changes of ν are dampened by the stochastic process interpolation.

As regards the number of free parameters to be estimated, the multi-period model yields estimates of $\nu_t^i,\ t = 1, \ldots,\ n,\ i = 1, \ldots,\ m$, for a total of nm free parameters. The drift vector μ and cholesky matrix A are then simply functions of the yearly parameter values. On the other hand, the HMM yields estimates of μ, A and ν_0 for a total of $\frac{m^2}{2} + \frac{5m}{2}$ free parameters. The yearly estimates of $\nu_t^i,\ t = 1, \ldots,\ n,\ i = 1, \ldots,\ m$, are then given as integrals, parameterized by μ, A and ν_0. Thus, the number of free parameters is lower for the HMM if $\frac{m^2}{2} + \frac{5m}{2} < nm$, or, equivalently, $m < 2n - 5$. For our sample of $n = 12$ years, the HMM has fewer free parameters if $m < 19$. Usually, we are content with a much lower dimension for the environment process ν.

Given these two major advantages, we conclude that, for the purpose of generating future scenarios, the HMM is preferred to the multi-period model. We will now try to refine the model by increasing the number of basis functions.

3.2 Three-factor model

Following [1], we implement the model from Section 2 with piecewise linear basis functions given by

$$\phi^1(x) = \begin{cases} 1 - \frac{x-25}{15} & \text{for } x \in [25, 40) \\ 0 & \text{for } x \in [40, 64], \end{cases}$$

$$\phi^2(x) = \begin{cases} \frac{1}{15}(x - 25) & \text{for } x \in [25, 40) \\ \frac{64}{24} - \frac{x}{24} & \text{for } x \in [40, 64], \end{cases}$$

$$\phi^3(x) = \begin{cases} 0 & \text{for } x \in [25, 40) \\ \frac{x}{24} - \frac{40}{24} & \text{for } x \in [40, 64]. \end{cases}$$

The linear combination $\sum_{i=1}^3 \nu_t^i \phi^i(x)$ is now piecewise linear and continuous, with mid point at $x = 40$ years. As in the two-factor model, the values of the factors are points on the logit inception probability curve: $\operatorname{logit} p_{25,t} = \nu_t^1$, $\operatorname{logit} p_{40,t} = \nu_t^2$ and $\operatorname{logit} p_{64,t} = \nu_t^3$.

The EM-algorithm stabilizes to within Monte Carlo error after about 250 iterations. The estimated inception probabilities are displayed in Figures 6-7. For reference, they are compared to the estimations from the multi-period model. Table 2 displays the estimated drift and volatility parameters from the HMM as a fraction of the corresponding estimates from the multi-period model.

Table 2: Relative difference of the estimated drift and volatility parameters between the two models.

	μ	σ
ν^1	0.45	0.32
ν^2	0.98	0.42
ν^3	0.93	0.29

The HMM seems to provide estimates of $\nu_{1:n}$ and μ that are quite close to the estimates from the multi-period model, at least for ν^2 and ν^3. For ν^1, the differences are quite pronounced from the year 2006 and onwards. It may be that the Brownian

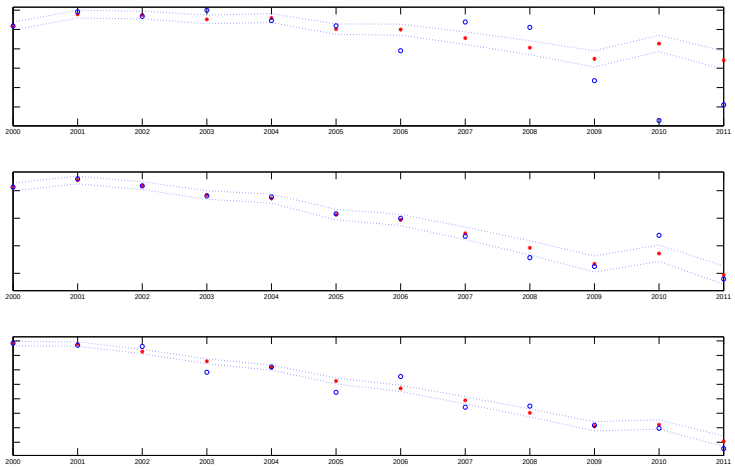

Figure 6: Estimates of $\nu_{1:n}$ (stars) with confidence bands (dashed). Estimates from [1] (circles) for comparison.

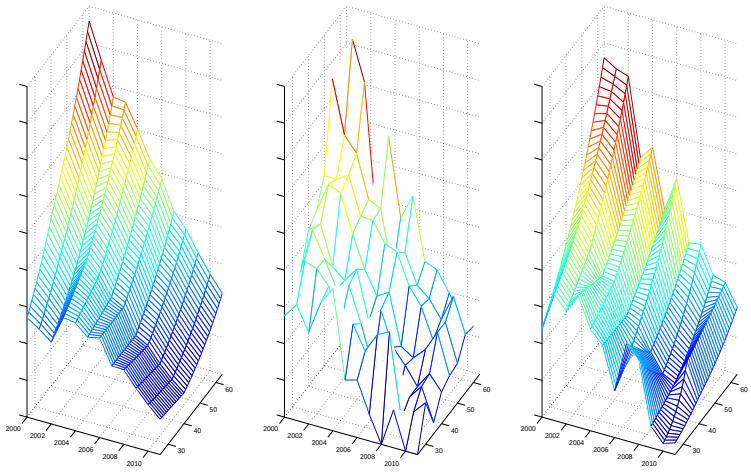

Figure 7: Left: Estimates of $p_{25:64,1:n}$. Middle: Raw data $D_{25:64,1:n}/E_{25:64,1:n}$. Right: Estimates from [1].

motion assumption could be altered, but this is a topic for future research. We return to this discussion point later on.

As was the case with the two-factor model, the estimated diffusion coefficients are considerably reduced. Judging from Figures 6-7, adding more complexity in terms of another underlying factor does not seem to significantly enhance the model. Here, the results are based on a midpoint of $x = 40$ years, but from extensive testing we find that the conclusion is valid for any choice of midpoint.

4 Disability termination model

The algorithm from Section 2 can easily be applied to obtain an extension of the disability termination model from [1].

Let $E_{x,d,t}$ be the number of individuals with disability inception age x and disability duration d at some point in the time period $[t, t+1)$. Further, let $R_{x,d,t}$ denote the number of individuals among $E_{x,d,t}$ with termination during $[t, t+1)$ and $[d, d+\Delta d)$. Further, let ν be a $k \times m$-dimensional Brownian motion starting at ν_0 with drift vector μ and Cholesky matrix A. We assume that the conditional distribution of $R_{x,d,t}$ given $E_{x,d,t}$ and ν_t is binomial:

$$R_{x,d,t} \sim \text{Bin}(E_{x,d,t}, p_{x,d,t}), \tag{29}$$

where

$$p_{x,d,t} := \frac{1}{1 + e^{-g(x,d,\nu_t)}} \tag{30}$$

denotes the probability that the disability of an individual, with disability inception age x and disability duration d at some point in the time period $[t, t+1)$, is terminated before duration $d + \Delta d$. Again, we adopt the basis function approach from [1], and choose a function $g : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^{k \times m} \mapsto \mathbb{R}$ of the form

$$g(x, d, \nu_t) = \sum_{i=1}^{m} \phi^i(x) \sum_{j=1}^{k} \psi^j(d) \nu_t^{i,j}, \tag{31}$$

where ϕ and ψ are basis functions in x and d, respectively.

Now, assume that we observe $R_{x,d,1:n} := (R_{x,d,1}, \ldots, R_{x,d,n})$, for x and d from given sets X and D of ages and disability durations, respectively. Let $\theta = (\mu, A, \nu_0)$. Then, the complete data log-likelihood is given by

$$\begin{aligned} l(\theta; R_{\cdot,\cdot,1:n}, \nu_{1:n}) = \sum_{t=1}^{n} \Big[\sum_{\substack{x \in X \\ d \in D}} \Big\{ R_{x,d,t} g(x, d, \nu_t) - E_{x,d,t} \log\big(1 + \exp\big\{g(x, d, \nu_t)\big\}\big) \Big\} \\ + \log f_{\nu_t | \nu_{t-1}}(\theta) + c_t \Big], \end{aligned} \tag{32}$$

where f denotes the density of ν_t given ν_{t-1}, and c_t is a constant. Integrating the log-likelihood and discarding terms that do not depend on θ, we again obtain

$$\begin{aligned} Q(\theta|\theta^k) = \sum_{t=1}^{n} \Big[-\frac{1}{2} E^{\theta^k} [(\nu_t - \nu_{t-1} - \mu)^T (AA^T)^{-1} (\nu_t - \nu_{t-1} - \mu) | R_{\cdot,\cdot,1:n}] \\ - \frac{1}{2} \log(\det(AA^T)) \Big]. \end{aligned} \tag{33}$$

The expectation and maximization steps are carried out exactly as in Sections 2.1-2.2.

5 Fitting Swedish termination rates

In this section, we implement the EM-algorithm from Sections 2.1-2.2 for the disability termination model from Section 4, and fit it to population data from Folksam.

5.1 Four-factor model

We propose an initial model for the termination probabilities with linear basis functions in both the age dimension x and the duration dimension d:

$$\begin{aligned}\phi^1(x) &= \frac{64-x}{39},\\ \phi^2(x) &= \frac{x-25}{39},\\ \psi^1(d) &= 1,\\ \psi^2(d) &= d.\end{aligned}$$

The obtained termination curve for a 25-year old is only affected by the term containing ϕ^1. Notice that, since $\phi^1(25) = 1$ and $\phi^2(25) = 0$, $\operatorname{logit} p(25, d, \nu_t) = \nu_t^{1,1} + d\nu_t^{1,2}$. In other words, the processes $\nu^{1,1}$ and $\nu^{1,2}$ correspond to the termination curve for a 25-year old. By the same argument, the processes $\nu^{2,1}$ and $\nu^{2,2}$ correspond to the termination curve for a 64-year old. The logistic conditional probability of termination for an x year old is thus a convex combination of the logistic conditional probabilities of termination for a 25-year old and a 64-year old.

The EM-algorithm stabilizes to within Monte Carlo error after about 260 iterations. The estimated parameters are displayed in Figure 8. Figures 9-10 displays the termination surface from the model alongside the corresponding Kaplan-Meier curves for the years 2006 and 2010, respectively. Results from [1] are displayed for reference. Table 3 displays the estimated drift and volatility parameters from the HMM as a fraction of the corresponding estimates from the multi-period model.

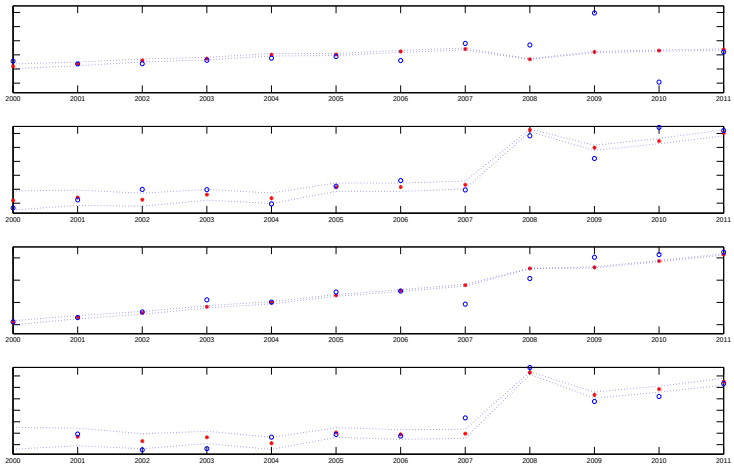

Figure 8: Estimates of $\nu_{1:n}$ (stars) with confidence bands (dashed). Estimates from [1] (circles) for comparison.

The HMM seems to provide estimates of $\nu_{1:n}$ and μ that are quite close to the estimates from the multi-period model, at least for ν^2, ν^3 and ν^4. For ν^1, the differences

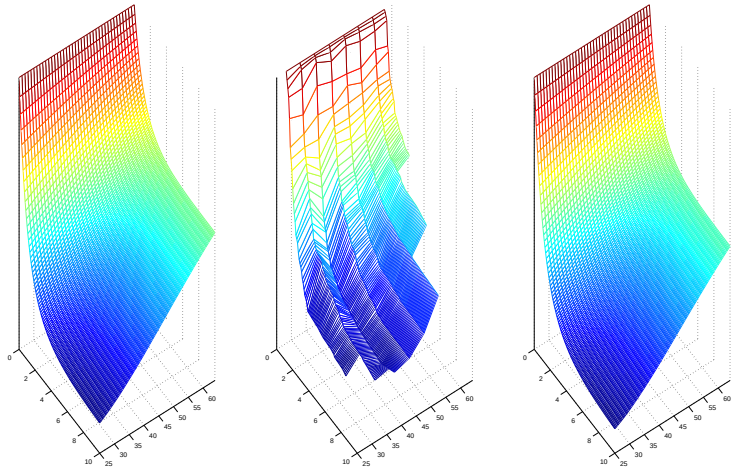

Figure 9: Left: Fitted termination surface, females 2006. Center: Kaplan-Meier. Right: Estimates from [1].

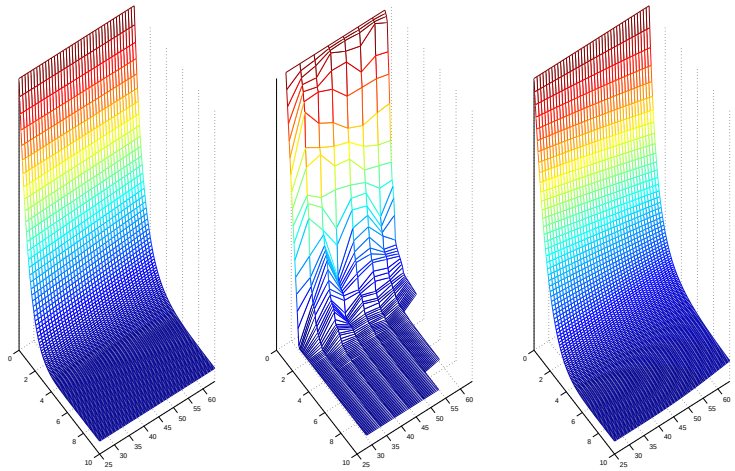

Figure 10: Left: Fitted termination surface, females 2010. Center: Kaplan-Meier. Right: Estimates from [1].

are quite pronounced from the year 2008 and onwards. Even so, the estimated termination surfaces from the HMM are close to the multi-period model surfaces for both 2006 and 2010. Further, we see that the estimated volatilities of $\nu^1, \ldots, \nu^4$ are lower for the HMM compared to the multi-period model, which is both reasonable and expected. The relatively large deviance for μ_1 can be explained by the fact that the corresponding estimate for the multi-period model is close to zero, so the absolute difference is small.

Table 3: Relative difference of the estimated drift and volatility parameters between the two models.

	μ	σ
ν^1	1.82	0.15
ν^2	0.88	0.75
ν^3	0.98	0.33
ν^4	1.08	0.90

5.2 Six-factor model

In order to refine the model we extend it to the following six-factor model:

$$\begin{aligned}
\phi^1(x) &= \frac{64-x}{39}, \\
\phi^2(x) &= \frac{x-25}{39}, \\
\psi^1(d) &= 1, \\
\psi^2(d) &= e^{-d}, \\
\psi^3(d) &= e^{-2d}.
\end{aligned}$$

The EM-algorithm stabilizes to within Monte Carlo error after about 140 iterations. The estimated parameters are displayed in Figure 11. Figures 12-13 display the termination surface from the model alongside the corresponding Kaplan-Meier curves for the years 2006 and 2010, respectively. Results from [1] are displayed for reference. Table 4 displays the estimated drift and volatility parameters from the HMM as a fraction of the corresponding estimates from the multi-period model.

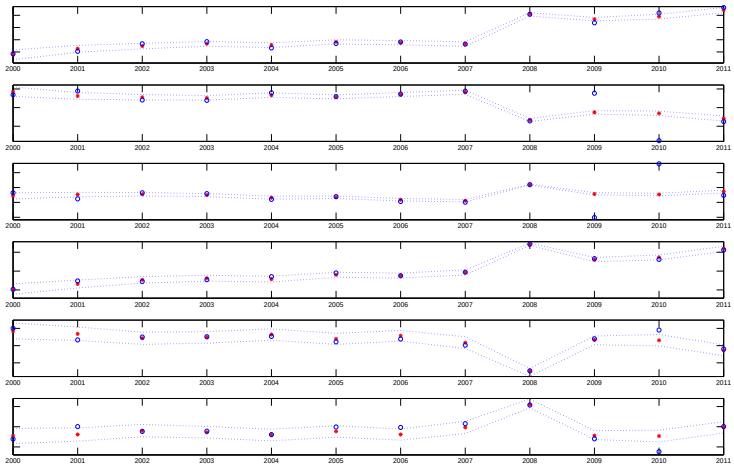

Figure 11: Estimates of $\nu_{1:n}$ (stars) with confidence bands (dashed). Estimates from [1] (circles) for comparison.

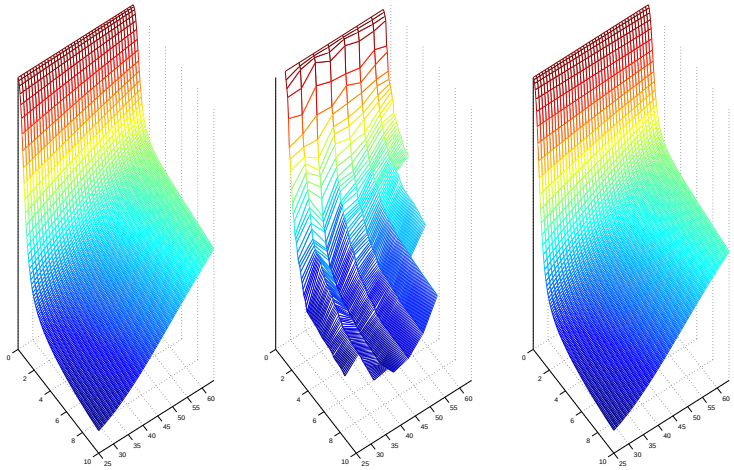

Figure 12: Left: Fitted termination surface, females 2006. Center: Kaplan-Meier. Right: Estimates from [1].

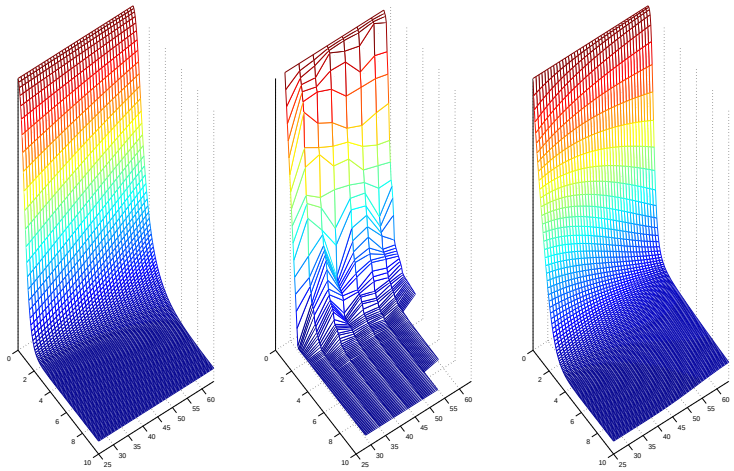

Figure 13: Left: Fitted termination surface, females 2010. Center: Kaplan-Meier. Right: Estimates from [1].

As for the four-factor model, we see that the estimated volatilities of $\nu^1, \ldots, \nu^6$ are lower for the HMM compared to the multi-period model. The estimates of $\nu_{1:n}$ and μ are similar to the estimates from the multi-period model, with some differences appearing from 2009 and onwards. Still, the estimated termination surfaces from the HMM are again close to the multi-period model surfaces for both 2006 and 2010, even though the estimated parameters for 2010 differ significantly between the two models. Since this was also the case with the four-factor model, it suggests a certain degree of

Table 4: Relative difference of the estimated drift and volatility parameters between the two models.

	μ	σ
ν^1	0.96	0.84
ν^2	0.99	0.41
ν^3	-1.56	0.24
ν^4	1.01	0.89
ν^5	0.91	0.85
ν^6	0.66	0.74

robustness for the fitting procedure. The relatively large deviance for μ_3, which even shows a change of sign, can be explained by the fact that the corresponding estimate for the multi-period model is close to zero, so that the absolute difference is small.

As a final note, we briefly comment on the nature of the unobservable environment process. The reform of the Swedish health insurance system in 2008 introduced harsher rules for obtaining benefits. On the other hand, for the proposed reform of 2014 the rules for obtaining benefits will become more lenient. Moreover, it is also possible that the population disability pattern follow the macro-economic trends of society in one way or another. All together, we are led to believe that the environment process follows a certain mean-reverting pattern. Unfortunately, the the data set used in this paper covers the rather short time period from 2000-2011, during which it may be hard to observe any mean-reverting pattern of the process. A topic for future research is to collect data from the new post-2014 regime and investigate whether mean-reverting processes such as the multivariate Vasicek model can be used to model the environment process.

6 Acknowledgements

The first author gratefully acknowledges financial support from the Swedish Export Credit Corp. (SEK). The second author gratefully acknowledges financial support from the Filip Lundberg and Eir's 50 Years foundations.

References

[1] ARO, H., DJEHICHE, B., AND LÖFDAHL, B. Stochastic modelling of disability insurance in a multi-period framework. *Scandinavian Actuarial Journal* (2013).

[2] ARO, H., AND PENNANEN, T. A user-friendly approach to stochastic mortality modelling. *European Actuarial Journal 1*, 2 (2011), 151–167.

[3] BROUHNS, N., DENUIT, M., AND VERMUNT, J. K. A poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics 31*, 3 (2002), 373–393.

[4] CHRISTIANSEN, M., DENUIT, M., AND LAZAR, D. The Solvency II square-root formula for systematic biometric risk. *Insurance: Mathematics and Economics 50* (2012), 257–265.

[5] CZADO, C., DELWARDE, A., AND DENUIT, M. Bayesian Poisson log-bilinear mortality projections. *Insurance: Mathematics and Economics 36*, 3 (2005), 260–284.

[6] DEMPSTER, A., LAIRD, N., AND RUBIN, D. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* (1977), 1–38.

[7] DJEHICHE, B., AND LÖFDAHL, B. Risk aggregation and stochastic claims reserving in disability insurance. *Insurance: Mathematics and Economics 59* (2014), 100–108.

[8] DUFFIE, D., ECKNER, A., HOREL, G., AND SAITA, L. Frailty correlated default. *The Journal of Finance 64*, 5 (2009), 2089–2123.

[9] FEYEL, D., AND ÜSTÜNEL, A. Log-concave measures. *TWMS Journal of Pure and Applied Mathematics 1*, 1 (2010), 92–105.

[10] HYNDMAN, R., AND ULLAH, M. S. Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics and Data Analysis 51* (2007), 4942–4956.

[11] LEE, R. D., AND CARTER, L. R. Modeling and forecasting us mortality. *Journal of the American statistical association 87*, 419 (1992), 659–671.

[12] RENSHAW, A., AND HABERMAN, S. Modelling recent time trends in uk permanent health insurance recovery, mortality and claim inception transition intensities. *Insurance: Mathematics and Economics 27* (2000), 365–396.

[13] WESTERBORN, J., AND OLSSON, J. Efficient particle-based online smoothing in general hidden Markov models. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on* (2014), IEEE, pp. 8003–8007.

[14] WU, C. J. On the convergence properties of the EM algorithm. *The Annals of statistics* (1983), 95–103.